arXiv:1611.07803v2 [cs.LO] 20 Jan 2017

Interpolation-Based GR(1) Assumptions Refinement*

Davide G. Cavezza and Dalal Alrajeh

Imperial College London, United Kingdom
{d.cavezza15,dalal.alrajeh}@imperial.ac.uk

Abstract. This paper considers the problem of assumptions refinement in the context of unrealizable specifications for reactive systems. We propose a new counterstrategy-guided synthesis approach for GR(1) specifications based on Craig's interpolants. Our interpolation-based method identifies causes for unrealizability and computes assumptions that directly target unrealizable cores, without the need for user input. Thereby, we discuss how this property reduces the maximum number of steps needed to converge to realizability compared with other techniques. We describe properties of interpolants that yield helpful assumptions in GR(1) and prove the soundness of the results. Finally, we demonstrate that our approach yields weaker assumptions than baseline techniques.

Keywords: Reactive synthesis; assumption refinement; interpolation

1 Introduction

Constructing formal specifications that capture user requirements precisely and from which implementations can be successfully derived is a difficult task [24]. Their imprecision often results from the conception of over-ideal systems, i.e., where the environment always behaves as expected [2,25]. Thus one of the challenges in building correct specifications is identifying sufficient assumptions over the environment under which a system would always be able to guarantee their satisfaction, in other words making a specification *realizable*.

This paper presents a new technique for automatically synthesizing assumptions over an adversarial environment for realizability assurance. More specifically, we develop a novel counterstrategy-guided synthesis procedure that iteratively generates assumption refinements, expressed in a fragment of Linear Temporal Logic (LTL) called *Generalized Reactivity (1)* (GR(1) for short), based on logical interpolation. *Craig interpolants* characterize automatically computable explanations for the inconsistency between Boolean formulae, in their shared alphabet. We exploit this feature to construct expressions that *explain* why a counterstrategy, and hence the environment, violates a guarantee, and whose negations form assumptions.

* To appear in TACAS 2017

We demonstrate in our case study applications that our approach *directly targets unrealizable cores*, in the sense that by adding the assumptions returned at each iteration, a specific subset of *minimally unfulfillable guarantees* [12] becomes realizable. Therefore each iteration takes a step closer to realizability. To characterize the scope of our approach we introduce the notion of fully-separable interpolants and prove the soundness of our computation when interpolants are fully separable. We further provide a discussion about the complexity of the proposed approach and its convergence, as well as the weakness of our refinements in comparison with those computed by existing techniques [26,3,4].

2 Related Work

Recent years have seen the development of effective counterstrategy-guided approaches to GR(1) assumptions refinement, notably [26,3,4]. Nonetheless those approaches depend significantly on users' knowledge of the problem domain and of the cause of unrealizability. The work in [26] requires users to specify a set of temporal logic templates as formulae with placeholders to be replaced with Boolean variables. Assumptions are then generated as instantiations of such templates that eliminate a given counterstrategy. This typically constrains the search space to only a subset of GR(1) formulae, which do not necessarily address the cause of unrealizability, and potentially eliminate viable solutions to the realizability problem. Similarly, the work in [3], although generating such templates automatically, requires users to provide a subset of variables for template instantiation. Unless the user knows the exact subset of variables that form the cause, this may yield assumptions that do not target the true cause of unrealizability, leading to refinements that needlessly over-constrain the environment. Our proposed method instead directly targets counter-strategies and unrealizable cores, and does not require users to provide variables for constructing refinements.

Other related work on assumption refinement includes those operating directly on game structures [11]. With regard to the parity game model used for controller synthesis (such as in [31,32]), the paper defines the concept of safety assumptions as sets of edges that have to be avoided by the environment, and the concept of fairness assumptions as sets of edges that have to be traversed by the environment infinitely often. The work devises an algorithm for finding minimal edge sets in order to ensure that the controller has a winning strategy. Our approach instead focuses on synthesizing general declarative temporal assertions whose inclusion has the effect of removing edges from the game structure, and directly targeting sources of unrealizability. The problem of synthesizing environment constraints has been tackled in the context of assume-guarantee reasoning for compositional model checking [30,14,20] to support compositional verification. In these, assumptions are typically expressed as LTSs and learning algorithms like $\mathcal{L}^*$ [5] are used to incrementally refine the environment assumptions needed in order to verify the satisfaction of properties.

Craig interpolants have been deployed in the context of abstraction refinement for verification in [18,17]. The differences with our work are in specification

language and overall objective: they seek additional assertions for static analysis of programs, while we look for GR(1) refinements of systems specifications to enable their automated synthesis. The authors of [16] use interpolation to support the extraction of pre- and trigger-conditions of operations within event-driven systems to enable the 'satisfaction' of goals expressed within restricted fragment of LTL. Though different in objective, approach and class of properties, our technique can help in identifying specifications operationalizable by [16].

3 Background

Generalized Reactivity (1) Specifications. LTL [28] is a formalism widely used for specifying reactive systems. The syntax of LTL is defined over a finite non-empty set of propositional variables $\mathcal{V}$, the logical constants *true* and *false*, Boolean connectives, and operators **X** (next), **G** (always), **F** (eventually), **U** (until). Given a set of states Q and a labelling function $\lambda : Q \to 2^{\mathcal{V}}$, an LTL formula ϕ is interpreted over an infinite sequence of states $\sigma = q_0 q_1 ...$ in the standard way, and its language $\mathcal{L}(\phi)$ is the set of (infinite) state valuation sequences $w = \lambda(q_0)\lambda(q_1)...$ such that $w \models \phi$). We assume that the set $\mathcal{V}$ consists of two disjoint sets: *input variables* $\mathcal{X}$ and *output variables* $\mathcal{Y}$. We will use the expression $B(\mathcal{V})$ for a boolean expression (i.e., a logical expression without temporal operators) which uses variables in the set $\mathcal{V}$. We will also denote by $\mathbf{X}\mathcal{V}$ the set of expressions obtained by prepending a "next" operator to the variables in $\mathcal{V}$: this is equivalent to the set of primed versions of such variables [9].

Generalized Reactivity (1) specifications (written GR(1) for short) are a subset of LTL of the form $\phi^{\mathcal{E}} \to \phi^{\mathcal{S}}$ where $\phi^{\mathcal{E}}$ represents the assumptions of an environment and $\phi^{\mathcal{S}}$ the guarantees of a controller. The expression ϕ^{θ}, where $\theta \in \{\mathcal{E}, \mathcal{S}\}$, is specified as conjunction of the following: (1) a Boolean formula φ^{θ}_{init} of the form $B(\mathcal{X})$ if $\theta = \mathcal{E}$ and $B(\mathcal{V})$ otherwise, representing initial conditions; (2) a set of LTL formulae φ^{θ}_{inv} of the form $\mathbf{G}B(\mathcal{V} \cup \mathbf{X}\mathcal{X})$ if $\theta = \mathcal{E}$ and $\mathbf{G}B(\mathcal{V} \cup \mathbf{X}\mathcal{V})$ when $\theta = \mathcal{S}$, representing invariants; and (3) a set of LTL formulae φ^{θ}_{fair} of the form $\mathbf{GF}B(\mathcal{V})$ representing fairness conditions. We will sometimes indicate GR(1) specifications as a tuple $\langle \phi^{\mathcal{E}}, \phi^{\mathcal{S}} \rangle$ with $\phi^{\theta} = \{\varphi^{\theta}_{init}\} \cup \{\varphi^{\theta}_{inv}\} \cup \{\varphi^{\theta}_{fair}\}$.

A finite-state Moore transducer is a tuple $M = \langle Q, q_0, \mathcal{I}, \mathcal{O}, \rho, \delta \rangle$ where Q is a set of states, $q_0 \in Q$ is the initial state, $\rho : Q \times \mathcal{I} \to Q$ is the transition function, and $\delta : Q \to \mathcal{O}$ is the output function. Given an input sequence $w = i_0 i_1 ...$, a run of M is the sequence $\sigma = q_0 q_1 ...$ such that $q_{k+1} = \rho(q_k, i_k)$ for all $k \geq 0$. A run σ on input sequence $w \in \mathcal{I}^{\omega}$ produces an infinite word $M(w) = (\delta(q_0), i_0), (\delta(q_1), i_1)....$ The language of a Moore transducer M is $\mathcal{L}(M) = \{M(w) | w \in \mathcal{I}^{\omega}\}$, i.e., the infinite words generated by a sequence of inputs and the corresponding outputs over runs of M. A Moore transducer M is said to satisfy an LTL expression ϕ if $\mathcal{L}(M) \subseteq \mathcal{L}(\phi)$; in this case we also say that M is a model of ϕ and we denote it as $M \models \phi$. A GR(1) property ϕ is said to be realizable if there exists an M (representing a controller) such that $M \models \phi$.

Given a specification $\langle \phi^{\mathcal{E}}, \phi^{\mathcal{S}} \rangle$ that is unrealizable, we say that $\varphi^{\mathcal{S}} \subseteq \phi^{\mathcal{S}}$ is *minimally unfulfillable w.r.t. to* $\phi^{\mathcal{E}}$ iff the removal of any guarantee $g \in \varphi^{\mathcal{S}}$

makes $\langle \phi^{\mathcal{E}}, \varphi^{\mathcal{S}} \backslash \{g\} \rangle$ realizable [12]. Furthermore, an assumption $a \in \phi^{\mathcal{E}}$ is said to be *unhelpful w.r.t.* $\phi^{\mathcal{S}}$ if $\forall \varphi^{\mathcal{S}} \subseteq \phi^{\mathcal{S}}$. $\langle \phi^{\mathcal{E}}, \varphi^{\mathcal{S}} \rangle$ is realizable $\leftrightarrow \langle \phi^{\mathcal{E}} \backslash \{a\}, \varphi^{\mathcal{S}} \rangle$ is realizable. It is said to be *helpful* otherwise. Given a set of minimally unfulfillable guarantees $\varphi^{\mathcal{S}}$ w.r.t. $\phi^{\mathcal{E}}$, let $\varphi^{\mathcal{E}} \subseteq \phi^{\mathcal{E}}$ be a set of helpful assumptions for $\varphi^{\mathcal{S}}$; the specification $\langle \varphi^{\mathcal{E}}, \varphi^{\mathcal{S}} \rangle$ is called an *unrealizable core* [12].

If a specification ϕ over $\mathcal{V} = \mathcal{X} \cup \mathcal{Y}$ is unrealizable, an unrealizable core $\langle \varphi^{\mathcal{E}}, \varphi^{\mathcal{S}} \rangle$ and an environment strategy (called a *counterstrategy*) can be computed [12,22]. A *counterstrategy* is defined as a Moore transducer $(S, s_{init}, 2^{\mathcal{Y}'}, 2^{\mathcal{X}}, \rho, \delta)$ that satisfies $\varphi^{\mathcal{E}}$ and violates $\varphi^{\mathcal{S}}$ [4]. It describes the inputs produced by an admissible environment in response to the output configuration yielded by the controller in order to force the violation of ϕ. The runs of a counterstrategy are called *plays*. The terms 'counterstrategy' and 'play' come from the game-theoretic algorithms used to reason about realizability [9,3,22]. The transition function ρ depends only on a subset of the output variables $\mathcal{Y}' \subseteq \mathcal{Y}$ [22]. We define a labelling function $\lambda' : S \rightarrow 2^{\mathcal{X} \cup \mathcal{Y}'}$ over states in the counterstrategy in this way: a propositional variable is in $\lambda'(s)$ if it is asserted in all the incoming transitions of s, while $\lambda'(s)$ is arbitrary for any s with no incoming transitions.

Interpolants. Craig interpolation was originally defined for first-order logic [15] and later for propositional logic [23]. No interpolation theorems have been proved for the general LTL. Extensions have been proposed recently for LTL fragments [21,19]. However these do not include GR(1) formulae and therefore are not applicable in our case. We use interpolation for propositional logic.

Formally, given an unsatisfiable conjunction of formulae $\alpha \wedge \beta$, a Craig interpolant I is a formula that is implied by α, is unsatisfiable in conjunction with β, and is defined on the common alphabet of α and β. We write $\mathcal{L}_\phi$ to denote the set of variables that occur in a formula ϕ (also called the *alphabet* of ϕ).

Definition 1 (Interpolant [23]). *Let α and β be two logical formulae such that their conjunction $\alpha \wedge \beta$ is unsatisfiable. Then there exists a third formula I, called* interpolant *of α and β, such that, $\alpha \rightarrow I$, $I \rightarrow \neg\beta$ and $\mathcal{L}_I \subseteq \mathcal{L}_\alpha \cap \mathcal{L}_\beta$.*

An interpolant can be considered as an over-approximation of α that is still unsatisfiable in conjunction with β. As stated in Craig's interpolation theorem, although an interpolant always exists, it is not unique. Several efficient algorithms have been proposed for interpolation in propositional logics. The resulting interpolant depends on the internal strategies of these algorithms (e.g., SAT solvers, theorem provers). Our approach is based on McMillan's interpolation algorithm described in [29] and implemented in MathSAT [13]. In brief, the algorithm considers a proof by resolution for the unsatisfiability of $\alpha \wedge \beta$. (See Appendix for details.)

4 Approach Overview

The general procedure is based on a sequence of realizability checks and counterstrategy computations, in the spirit of [26,3]. A specification $\langle \phi^{\mathcal{E}}, \phi^{\mathcal{S}} \rangle$ is first

checked for realizability. If it is unrealizable, a counterstrategy C and an unrealizable core $\langle \varphi^{\mathcal{E}}, \varphi^{\mathcal{S}} \rangle$ are computed. The counterstrategy constitutes an example of environment behaviours that force the violation of the guarantees of $\varphi^{\mathcal{S}}$: therefore, the assumptions $\varphi^{\mathcal{E}}$ are refined by adding a GR(1) formula which is inconsistent with the counterstrategy. A set of such formulae Ψ is automatically computed by interpolating (α) the description of an environment behaviour in the counterstrategy, given by the assumptions and a sequence of state labellings in the counterstrategy; and (β) the guarantees, and by negating the interpolant. A formula $\psi_i \in \Psi$ is added to the original set of assumptions $\phi^{\mathcal{E}}$ and the procedure repeats the above steps recursively until realizability is achieved. Algorithm 1 describes this procedure schematically.

Algorithm 1: CounterstrategyGuidedRefinement procedure

Data: $\phi^{\mathcal{E}}$, assumptions
Data: $\phi^{\mathcal{S}}$, guarantees
Result: $\{\psi_i\}$, set of alternative assumption refinements such that $\phi^{\mathcal{E}} \wedge \psi_i \rightarrow \phi^{\mathcal{S}}$ is realizable for every i

```
if Satisfiable(φ^E → φ^S) & not Realizable(φ^E → φ^S) then
    (φ^E, φ^S, C) := Counterstrategy(φ^E, φ^S);
    Ψ := InterpolationBasedSynthesis(φ^E, φ^S, C);
    foreach ψ_i ∈ Ψ do
        foreach ψ'_j ∈ CounterstrategyGuidedRefinement(φ^E ∧ ψ_i, φ^S) do
            refinements.add(ψ_i ∧ ψ'_j) ;
        end
    end
    return refinements;
else if Satisfiable(φ^E → φ^S) & Realizable(φ^E → φ^S) then
    return {true};
else
    return {false};
```

The function **InterpolationBasedSynthesis** constitutes the core of our proposal (see Algorithm 2). It takes as inputs an unrealizable core and a counterstrategy and executes the computation of Ψ via interpolation. We give the details in the following section.

5 Interpolation-Based Synthesis

Each execution of **InterpolationBasedSynthesis** involves extracting temporal formulae that are satisfied by a single play of a counterstrategy (henceforth called *counterplay*), and obtaining refinements from its negation. It is sufficient to exclude a single counterplay of a counterstrategy to eliminate the entire counterstrategy from models of the assumption. Reasoning about counterplays has also some advantages, which are discussed in Sect. 8. For the purpose of this

Algorithm 2: InterpolationBasedSynthesis($\varphi^{\mathcal{E}}$, $\varphi^{\mathcal{S}}$, C)

Data: $\varphi^{\mathcal{E}}$, environment assumptions (in an unrealizable core)
Data: $\varphi^{\mathcal{S}}$, system guarantees (in an unrealizable core)
Data: C, counterstrategy
Result: Ψ, alternative assumptions eliminating the counterstrategy

```
1  π_C := ExtractCounterplay(C);
2  u := 0;
3  π_{C,u} := π_C;
4  Ψ_old := ∅;
5  stopping_condition := true;
6  repeat
7  |  [[π_{C,u}, φ^E_u]]:= TranslateCounterplayAssumptions(π_{C,u}, φ^E);
8  |  [[φ^S_u]] := TranslateGuarantees(π_{C,u},φ^S);
9  |  I_u := Interpolate([[π_{C,u}, φ^E_u]], [[φ^S_u]]);
10 |  if I_u == false or I_u is not fully-separable then
11 |  |  Ψ := false;
12 |  |  stopping_condition := true;
13 |  else
14 |  |  T(I_u) := TranslateInterpolant(π_{C,u}, I_u);
15 |  |  Ψ := ExtractDisjuncts(¬T(I_u));
16 |  |  if π_{C,u} is looping then
17 |  |  |  if Ψ ≠ Ψ_old then
18 |  |  |  |  Ψ_old := Ψ;
19 |  |  |  |  u := u + 1;
20 |  |  |  |  π_{C,u} := UnrollCounterplay(π_C,u);
21 |  |  |  |  stopping_condition:= false;
22 |  |  |  else
23 |  |  |  |  stopping_condition := true;
24 |  |  |  end
25 |  |  end
26 |  end
27 until stopping_condition;
28 return Ψ;
```

paper, we assume that the procedure **ExtractCounterplay** (line 1) extracts a counterplay π_C at random and consider metrics for selecting one in future work. A counterplay representing the violation of an initial condition or an invariant guarantee is finite, while that of a fairness guarantee violation ends in a loop [27]. We call the latter a *looping counterplay*, and the loop an *ending loop*.

We distinguish four types of states that may appear in π_C: (a) the *initial state* $S^{init} = \{s^{init}\}$; (b) the *failing state* in a finite counterplay $S^{fail} = \{s^{fail}\}$ (c) *looping states* that include the states in ending loop, $S^{loop} = \{s_1^{loop}, \ldots, s_h^{loop}\}$, (d) *transient states* including all states between the initial state and the first failing state or loop state (exclusive) $S^{trans} = \{s_1^{trans}, \ldots, s_k^{trans}\}$. With this classification, a finite counterplay has the form $s^{init} s_1^{trans} \ldots s_k^{trans} s^{fail}$; whilst

a looping counterplay has the form $s^{init}s_1^{trans}\dots s_k^{trans}(s_1^{loop}\dots s_h^{loop})^\omega$. The formulae in the next subsection also refer to a fifth set of states, called *unrolled states*, which represent replicates of looping states, and to the unrolling degree u. They are explored in Sect. 5.2. Each state in π_C is labelled with variables from the set $\mathcal{X}\cup\mathcal{Y}'$ defined in Sect. 3. The value of u is initialized to 0, and thus $\pi_{C,u}$ equates to π_C (lines 2–3).

The extraction of the counterplay occurs at the start of every call of the synthesis phase. The remaining steps described in this section are iteratively executed when the extracted counterplay is looping, and only once otherwise. In the former case, we will refer to the iteration as the *inner-cycle*, to distinguish it from the counterstrategy-guided refinement cycle.

5.1 Candidate Assumptions Computation

Refinements of environment assumptions are computed in four steps: (*i*) production of two inconsistent Boolean formulae from the counterplay and the unrealizable core, (*ii*) interpolation between the two Boolean formulae, (*iii*) translation of the interpolant into LTL, and (*iv*) negation of the translated interpolant.

Step (*i*) is executed by the functions **TranslateCounterplayAssumptions** and **TranslateGuarantees** (lines 7-8). The procedure employs the translation scheme in [7] for bounded model checking: it ensures that the obtained Boolean formula is satisfiable if and only if the play taken into account satisfies the LTL formula. The inclusion of assumptions in the counterplay translation is important in yielding an interpolant in the shared alphabet of assumptions and guarantees that explains why the assumptions violate the guarantees. Given a GR(1) formula in φ^θ over $\mathcal{V}$ and a counterplay $\pi_{C,u}$ with state space $S_\pi\in S$, its translation is a Boolean formula over the domain $\mathcal{V}(S_\pi)$ obtained by replicating every variable $p\in\mathcal{V}$ for every state $s\in S_\pi$; we denote by $p(s)$ the replica of p corresponding to state s, and by $\mathcal{V}(s)$ the subset of $\mathcal{V}(S_\pi)$ containing all the variables referring to state s. This step produces two formulae:

- $[[\pi_{C,u},\varphi_u^{\mathcal{E}}]]$, which is a conjunction between the assumptions translation $[[\varphi_u^{\mathcal{E}}]]$ over $\pi_{C,u}$ and a formula representing the valuation of every $s\in S_\pi$ in $\pi_{C,u}$; the latter is a conjunctive formula containing a literal $p(s)$ (resp. $\neg p(s)$) for every $p\in\mathcal{X}\cup\mathcal{Y}'$ that is true (resp. false) in $\lambda'(s)$ (see end of Sect. 3);
- $[[\varphi_u^{\mathcal{S}}]]$ which is the guarantees translation over $\pi_{C,u}$.

(The translations $[[\varphi_u^{\mathcal{E}}]]$ and $[[\varphi_u^{\mathcal{S}}]]$ are given in the Appendix.) Since by definition a counterplay π_C satisfies the assumptions and violates the guarantees, the formula $[[\pi_{C,u},\varphi_u^{\mathcal{E}}]]\wedge[[\varphi_u^{\mathcal{S}}]]$ is unsatisfiable by construction. Therefore, there exists an interpolant for $[[\pi_{C,u},\varphi_u^{\mathcal{E}}]]$ and $[[\varphi_u^{\mathcal{S}}]]$.

Step (*ii*) consists of the function **Interpolate** (line 9). The returned interpolant I_u is an over-approximation of $[[\pi_{C,u},\varphi_u^{\mathcal{E}}]]$ which by definition implies the negation of $[[\varphi_u^{\mathcal{S}}]]$: it can be interpreted as a cause of the guarantees not being satisfied by the counterplay, and as such a characterization of a set of counterplays not satisfying the guarantees.

From such interpolant the procedure aims at extracting a set of refinements that fit the GR(1) format. In order to do this, the Boolean to temporal translation requires the interpolant to adhere a specific structure. This is embodied in the notion of *full-separability*. To formally define full-separability, we need first to define state-separability and I/O-separability.

Definition 2 (State-separable interpolant). *An interpolant I_u is said to be* state-separable *iff it can be expressed as*

$$\bigwedge_{s \in S_u} B_s(\mathcal{V}(s)) \tag{1}$$

where $B_s(\mathcal{V}(s))$ is a Boolean formula either equal to true or expressed over variables in $\mathcal{V}(s)$ only.

We will refer to each $B_s(\mathcal{V}(s))$ as a *state component* of the interpolant. In particular, a state component is equal to *true* if I_u does not use any variables from s. State-separability intuitively means that the subformulae of the interpolant involving a single state are linked by conjunctions. This means that in any model of the interpolant each state component must be itself *true*.

Definition 3 (I/O-separable Boolean expression). *A Boolean expression $B_s(\mathcal{V}(s))$ is said to be* I/O-separable *if it can be written as a conjunction of two subformulae containing only input and output variables respectively:*

$$B_s(\mathcal{V}(s)) = B_{s,\mathcal{X}}(\mathcal{X}(s)) \wedge B_{s,\mathcal{Y}}(\mathcal{Y}(s)) \tag{2}$$

We call $B_{s,\mathcal{X}}(\mathcal{X}(s))$ and $B_{s,\mathcal{Y}}(\mathcal{Y}(s))$ the *projections* of $B_s(\mathcal{V}(s))$ onto $\mathcal{X}$ and $\mathcal{Y}$ respectively. Any model of an I/O-separable Boolean expression satisfies the projections separately. We can now define full-separability of an interpolant.

Definition 4 (Fully-separable interpolant). *An interpolant is called* fully-separable *if it is state-separable and each of its state components is I/O-separable.*

An example of a fully-separable interpolant over $\mathcal{X} = \{a, b\}, \mathcal{Y} = \{c, d\}$ and states $S = \{s_0, s_1\}$ is $(a(s_0) \vee b(s_0)) \wedge c(s_0) \wedge \neg b(s_1)$; a non-fully-separable interpolant, instead, is $a(s_0) \vee a(s_1)$, since literals referring to different states are linked via a disjunction.

Remark 1. A particular class of fully-separable interpolants is that of fully conjunctive interpolants, where no disjunctions appear. Whether or not the resulting interpolant is conjunctive depends on the order in which the interpolation algorithm [29] chooses the root clauses for building the unsatisfiability proof. A sufficient condition for obtaining a fully-conjunctive interpolant is that such root clauses be single literals from $[[\pi_{C,u}, \varphi_u^{\mathcal{E}}]]$, and that the pivot variable in each resolution step belong to the shared alphabet of $[[\pi_{C,u}, \varphi_u^{\mathcal{E}}]]$ and $[[\varphi_u^{\mathcal{S}}]]$. (see Appendix for details on the interpolation algorithm used).

Step (*iii*) consists of the function **TranslateInterpolant** (line 14). It converts a fully-separable interpolant $I_u = \bigwedge_{s \in S_u} B_s(\mathcal{V}(s))$ into the LTL formula

$$\begin{aligned} \mathcal{T}(I_u) = & B^{init}_{\mathcal{X}}(\mathcal{X}) \wedge \bigwedge_{s \in S_u} \mathbf{F}\left(B_s(\mathcal{V}) \wedge B_{\mathrm{succ}\,(s),\mathcal{X}}(\mathbf{X}\mathcal{X})\right) \wedge \\ & \mathbf{FG} \bigvee_{j=1}^{|S^{loop}|} \left(B^{loop}_j(\mathcal{V}) \wedge \bigwedge_{r=1}^{u} B^{unr}_{j,r}(\mathcal{V})\right) \end{aligned} \tag{3}$$

where the expression $B^{init}_{\mathcal{X}}(\mathcal{X})$ is a shorthand for $B_{s^{init},\mathcal{X}}(\mathcal{X})$, $B^{loop}_j(\mathcal{V})$ for $B_{s^{loop}_j}(\mathcal{V})$ and $B^{unr}_{j,r}(\mathcal{V})$ for $B_{s^{unr}_{j,r}}(\mathcal{V})$. Formula (3) is formed from the single state components of I_u by replacing the variables in $\mathcal{V}(s)$ with the corresponding variables in $\mathcal{V}$ and by projecting the components onto the input variables where required by the GR(1) template. The translation consists of three units: a subformula describing the initial state, a conjunction of **F** formulae each containing two consecutive state components, and an **FG** formula; this unit consists of a disjunction over all the looping states, where each disjunct j groups the state components of all the replicas of state s^{loop}_j.

Formula (3) is guaranteed to hold in the counterplay π_C. Intuitively, since I_u is fully-separable by construction, $[[\pi_{C,u}, \varphi^{\mathcal{E}}_u]]$ implies each state component and its projections onto $\mathcal{X}$ and $\mathcal{Y}'$. A state component $B_s(\mathcal{V}(s))$ corresponds to a formula $B_s(\mathcal{V})$ satisfied by state s of the counterplay. Therefore, since the initial state satisfies $B^{init}(\mathcal{V})$, π_C satisfies $B^{init}_{\mathcal{X}}(\mathcal{X})$; since there are two consecutive states s and $\mathrm{succ}\,(s)$ that satisfy $B_s(\mathcal{V}(s))$ and $B_s(\mathcal{V}(\mathrm{succ}\,(s)))$ respectively, π_C satisfies $\mathbf{F}\left(B_s(\mathcal{V}) \wedge B_{\mathrm{succ}\,(s),\mathcal{X}}(\mathbf{X}\mathcal{X})\right)$. Finally, for the **FG** subformula, it is sufficient to observe that the looping state j satisfies the formula obtained from the state components referring to s^{loop}_j and $s^{unr}_{j,r}$: since the counterplay remains indefinitely in the looping state, there is a suffix of it where such formula is true for at least one j. Based on these considerations, we prove the following soundness property.

Theorem 1. *Let π_C be a counterplay and $\varphi^{\mathcal{E}}$ a set of assumptions satisfied in π_C, such that their Boolean translation $[[\pi_{C,u}, \varphi^{\mathcal{E}}_u]]$ implies I_u, and let I_u be a fully-separable interpolant. Then $\pi_C \models \mathcal{T}(I_u)$.*

The proof is in the Appendix. In the case a fully-separable interpolant is not generated from which $\mathcal{T}(I_u)$ can be constructed, the algorithm returns *false* as its candidate assumption. Otherwise, the approach proceeds to step (*iv*) (function **ExtractDisjuncts**, line 15) producing the candidate refinements by negating (3) and extracting the disjuncts in the resulting formula:

$$\begin{aligned} \neg B^{init}_{\mathcal{X}}(\mathcal{X}) \vee & \bigvee_{s \in S_u} \mathbf{G}\neg\left(B_s(\mathcal{V}) \wedge B_{\mathrm{succ}\,(s),\mathcal{X}}(\mathbf{X}\mathcal{X})\right) \vee \\ & \mathbf{GF} \bigwedge_{j=1}^{|S^{loop}|} \neg\left(B^{loop}_j(\mathcal{V}) \wedge \bigwedge_{r=1}^{u} B^{unr}_{j,r}(\mathcal{V})\right) \end{aligned} \tag{4}$$

Each disjunct above is a GR(1) candidate assumption which, by Theorem 1, ensures the exclusion of the counterplay π_C from the models of the assumptions.

5.2 Equivalence Checking and Unrolling

The equivalence checking of the produced candidates and the unrolling of the counterplay (lines 17-24) are only executed in case of a looping counterplay. Thus in each iteration of the inner-cycle, our procedure checks whether the synthesized assumptions are equivalent to the assumptions $\boldsymbol{\psi}_{old}$ computed in the previous iteration. If not, the looping part of the counterplay is unrolled once (**UnrollCounterplay**, line 20) and the steps in Sect. 5.1–5.2 are repeated. If the equivalence condition is met, the synthesis procedure returns the last set of computed candidates as output.

Counterplay unrolling consists in making the first traversals of looping states explicit. It is achieved by augmenting a counterplay with replicates of the looping states. The number of unrollings is referred to as the *unrolling degree* u. Each unrolling yields a new set of states $S^{unr} = \{s^{unr}_{1,1}, \ldots, s^{unr}_{h,1}, \ldots, s^{unr}_{1,u}, \ldots, s^{unr}_{h,u}\}$. An unrolled looping counterplay has the form $s^{init} s^{trans}_1 \ldots s^{trans}_k s^{unr}_{1,1} \ldots s^{unr}_{h,1} \ldots s^{unr}_{1,u} \ldots s^{unr}_{h,u} (s^{loop}_1 \ldots s^{loop}_k)^\omega$. Unrolling has two possible effects on the computed interpolant: on one hand, it can introduce new state components in the interpolant, which yield new invariant refinements according to (4); on the other hand, the interpolant can express a more specific characterization of looping states, which corresponds to a weaker fairness refinement in (4). These effects are both observed in our evaluation (see Sect. 7).

6 Convergence

Our procedure is guaranteed to terminate after a finite number of recursive calls. We discuss below the case of all computed interpolants being fully-separable. If not, the procedure terminates with a trivial assumption refinement *false*.

Theorem 2. *Given a satisfiable but unrealizable specification* $\langle \phi^{\mathcal{E}}, \phi^{\mathcal{S}} \rangle$ *Algorithm 1 terminates with a realizable specification* $\langle \phi^{\mathcal{E}'}, \phi^{\mathcal{S}} \rangle$.

To prove this, it is sufficient to show that both the recursion in Algorithm 1 and the iteration over unrollings in Algorithm 2 reach the respective termination conditions. In the following arguments, we will refer to the recursion tree of Algorithm 1. Each node is associated with the candidate assumption tested in one specific call of **CounterstrategyGuidedRefinements**. The root corresponds to the initial assumption; every internal node symbolizes an unrealizable assumptions refinement; the children of an internal node correspond to the alternative refinements that rule out the relevant counterstrategy. The leaves represent alternative realizable assumption refinements returned by the algorithm. We will show that this tree has finite depth and breadth.

Let us consider the number of children n_C of an internal node (the subscript C indicates the counterstrategy computed in that internal node). It consists of

the maximum number of refinements that are generated from a single counterstrategy. Assuming that the maximum unrolling degree is finite (we will see that later in this section), denoted $u_{C,MAX}$, the maximum number of refinements generated from C can be computed by counting the maximum number of disjuncts in (4). Suppose $|S_{u_{C,MAX}}|$ denotes the number of distinct states in the unrolled counterplay, then $n_C \leq |S_{u_{C,MAX}}| + 2$: we count one initial condition, one fairness condition and $|S_{u_{C,MAX}}|$ invariants. Given that every node has a finite number of children, the breadth of each level in the tree is also finite.

We now consider the depth. The algorithm keeps refining a computed assumption until the property becomes realizable (in case the returned refinement is *false*, then the property is realizable, and therefore the algorithm reaches a *true* leaf). Given the soundness property, at each step every refinement excludes the latest computed counterstrategy; since this counterstrategy satisfies all the previously computed refinements by definition, the new refinement cannot be equivalent to any of the previous refinements along the same branch.

For the above reason, the depth d of the recursion tree is limited by the maximum number of existing GR(1) refinements modulo logical equivalence. The maximum number of initial conditions is $d_{init,MAX} = 2^{2^{|\mathcal{X}|}}$, that is the number of all distinct Boolean expressions over the input variables. The maximum number of invariants is $d_{inv,MAX} = 2^{2^{|\mathcal{V}|}+2^{|\mathcal{X}|}}$; this corresponds to the maximum number of distinct B_s that can be present in the expression (4) times the number of distinct $B_{\mathrm{succ}\,(s),X}$. Finally, the maximum number of distinct fairness assumptions is $d_{fair,MAX} = 2^{2^{|\mathcal{V}|}}$ Therefore, the total depth d is bounded by the sum of these three quantities: $d \leq d_{MAX} = d_{init,MAX} + d_{inv,MAX} + d_{fair,MAX}$.

Given the above, we conclude that the recursion tree is finite. This gives us a worst-case upper bound on the depth d of the recursion, which has a doubly exponential growth over $|\mathcal{V}|$ — a general observation of counterstrategy-guided assumptions refinement strategies. It remains to show that the inner-cycle terminates in finite time. As mentioned in Sect. 5.2, each iteration can provide additional or weaker refinements with respect to the previous iteration. The termination condition holds when the current iteration does not yield new refinements with respect to the previous one. This is reached in the worst case after all distinct GR(1) refinements are generated. The computation is the same as the one for d: $u_{C,MAX} = d_{MAX}$.

7 Evaluation

We apply our approach to two benchmarks presented in [3,9,22]: a lift controller and ARM's AMBA-AHB protocol. The requirements analysis tool RATSY [8] is used to check unrealizability and compute counterstrategies. The SAT solver MathSAT [13,10] is used to compute interpolants. We implemented a translation module for GR(1) specifications and randomly extracted plays into a propositional logic format executable by MathSAT. For each case study, we report the maximum depth and breadth of the recursion tree, and an interpretation of some interesting refinements that are computed. Details are available at [1].

Table 1 provides a summary of both case studies. The columns **In** and **Out** contain the number of input and output variables in the specification alphabet respectively; **A** and **G** contain the number of assumptions and guarantees respectively; **MaxPlay** contains the maximum number of states in a counterplay among all the counterplays used in the refinement process; **MaxUnr** reports the maximum unrolling degree reached in any step of the approach before reaching the termination condition; **TreeDepth** corresponds to the depth of the recursion tree; **MaxAltRef** is the maximum number of alternative refinements computed to rule out any single counterstrategy (it corresponds to the maximum number of children of an internal node in the recursion tree); **#Ref** shows the total number of refinement sets computed that make the property realizable.

Table 1. Summary of refinement results on benchmarks

Specification	In	Out	A	G	MaxPlay	MaxUnr	TreeDepth	MaxAltRef	#Ref
Lift	3	3	7	12	2	2	1	3	3
AMBA02	7	16	10	66	4	2	3	6	17
AMBA04	11	23	16	97	7	1	2	2	8
AMBA08	19	36	28	157	18	1	7	2	80

7.1 Lift Controller

This case study (also used for controller synthesis problems [9,3]) involves the specification of a system comprising a lift controller. The lift moves between three floors. The environment consists of three buttons, whose states can be *pressed* or *unpressed*; the corresponding state is represented by three binary input variables $\{b_1, b_2, b_3\}$. The controller's state consists of three output variables $\{f_1, f_2, f_3\}$ which indicate at which floor the lift is. The assumptions are:

1. $\varphi^e_{init} = \neg b_1 \wedge \neg b_2 \wedge \neg b_3$
2. $\varphi^e_{1,i} = \mathbf{G}(b_i \wedge f_i \rightarrow \mathbf{X}\neg b_i)$
3. $\varphi^e_{2,i} = \mathbf{G}(b_i \wedge \neg f_i \rightarrow \mathbf{X} b_i)$

for $i \in \{1, 2, 3\}$. They state that the buttons are not pressed in the initial state (1); a pressed button transits to a non-pressed state when the lift arrives at the corresponding floor (2); and the button remains in the pressed state until the lift arrives at that floor (3). The guarantees are:

1. $\varphi^s_{init} = f_1 \wedge \neg f_2 \wedge \neg f_3$
2. $\varphi^s_1 = \mathbf{G}(\neg(f_1 \wedge f_2) \wedge \neg(f_2 \wedge f_3) \wedge \neg(f_1 \wedge f_3))$
3. $\varphi^s_{2,1} = \mathbf{G}(f_1 \rightarrow (\mathbf{X} f_1 \vee \mathbf{X} f_2))$
4. $\varphi^s_{2,2} = \mathbf{G}(f_2 \rightarrow (\mathbf{X} f_1 \vee \mathbf{X} f_2 \vee \mathbf{X} f_3))$
5. $\varphi^s_{2,3} = \mathbf{G}(f_3 \rightarrow (\mathbf{X} f_2 \vee \mathbf{X} f_3))$
6. $\varphi^s_3 = \mathbf{G}(((f_1 \wedge \mathbf{X} f_2) \vee (f_2 \wedge \mathbf{X} f_3) \vee (f_2 \wedge \mathbf{X} f_1) \vee (f_3 \wedge \mathbf{X} f_2)) \rightarrow (b_1 \vee b_2 \vee b_3))$
7. $\varphi^s_{4,i} = \mathbf{GF}(b_i \rightarrow f_i)$
8. $\varphi^s_{5,i} = \mathbf{GF} f_i$

for $i \in \{1, 2, 3\}$. They state that the lift starts from floor 1 (1); it can never be in two floors at the same time (2); it can move only between consecutive states (3-5), and moves only when at least a button is pressed (6); plays in which the environment keeps a button b_i pressed infinitely and the lift never reaches the corresponding f_i are forbidden (7); and that the lift is required to visit all the floors infinitely often (8). Given this specification, the fairness guarantee can be satisfied if the environment sets one of its b_i to 1 at least once.

The specification is unrealizable, since when the buttons (environment) stay indefinitely unpressed, the lift (controller) cannot move and therefore $\varphi^s_{5,2}$ and $\varphi^s_{5,3}$ are violated. The unrealizable core consists of the whole set of assumptions and the guarantees φ^s_{init}, $\varphi^s_{2,1}$, φ^s_3 and $\varphi^s_{5,2}$. From this core, RATSY computes the counterstrategy π_C in Fig. 1, which consists of a unique play. After translating the unrealizable core over the counterplay, the interpolant is $I_0 = \neg b_1(s_0) \wedge \neg b_2(s_0) \wedge \neg b_3(s_0)$, which corresponds to the GR(1) refinement $\neg b_1 \wedge \neg b_2 \wedge \neg b_3$. The first unrolling is performed yielding the interpolant $I_1 = \neg b_1(s_0) \wedge \neg b_2(s_0) \wedge \neg b_3(s_0) \wedge \neg b_1(s^{unr}_{1,1}) \wedge \neg b_2(s^{unr}_{1,1}) \wedge \neg b_3(s^{unr}_{1,1})$. By translating and negating this interpolant, we obtain the alternative refinements

1. $b_1 \vee b_2 \vee b_3$
2. $\mathbf{G}(\neg b_1 \wedge \neg b_2 \wedge \neg b_3 \rightarrow \mathbf{X}(b_1 \vee b_2 \vee b_3))$
3. $\mathbf{GF}(b_1 \vee b_2 \vee b_3)$

Notice that unrolling results in an interpolant containing an additional state component, thus allowing for more alternative refinements (see Sect. 5.2). Moreover, the new state component refers to an unrolled state, from which a new fairness refinement that is not inferable from I_0 is synthesized. The second unrolling produces equivalent refinements, and thereby the inner-cycle terminates.

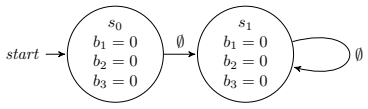

Fig. 1. Lift counterstrategy produced by RATSY. The labelling λ' is shown in each state. In this case the lift position plays no role in the environment's choice of next state, therefore $\mathcal{Y}' = \emptyset$.

Every candidate refinement computed by our approach is helpful. Moreover, each one solves the unrealizability problem for the original specification. Refinement (1) does this in a trivial way, since it contradicts the initial assumption contained in the specification. Notice that all the computed refinements force at least one of the buttons to be pressed at some point in any play of the environment. This corresponds to the refinement produced by the approach in [3].

7.2 AMBA-AHB Protocol

The *Advanced High-performance Bus* (*AHB*) is part of the *Advanced Microcontroller Bus Architecture* (*AMBA*) specification. It is an open-source communication protocol for on-chip devices through a shared bus. Devices are divided into

masters, which initiate a communication, and *slaves*, which respond to requests. Multiple masters can request the bus simultaneously, but only one at a time can communicate through it. Masters and slaves constitute the environment, while the system is the bus arbiter implementing the protocol. The specification of the AHB protocol is provided with RATSY. It is a GR(1) description of the protocol in [6], and formalized in [9]. We consider specifications for two, four and eight masters (AMBA02, AMBA04, AMBA08 respectively) which are realizable. To evaluate our approach, we remove the assumption $\mathbf{GF}$*hready* as done in [3,26].

In all the variants, our approach was able to produce refinements that were semantically related to the removed assumption. In the AMBA02 case, one of the refinements is the invariant $\mathbf{G}(\mathit{hready} \vee \mathbf{X}\mathit{hready})$, which forces *hready* to be true at least every two steps. The other refinements in all the AMBA0x variants involve the variable *hmaster*, which indicates the master that currently owns the bus. These refinements force *hmaster* to change infinitely often. This corresponds to having *hready* equal to *true* infinitely often, since *hready* must be true at any ownership switch according to the protocol [9].

The approach was further tested by extracting different counterplays from the same counterstrategy in the AMBA02 case. Every refinement produced within each synthesis call was helpful. We compared our results to those obtainable through [3,26] when the set of shared alphabet is provided as input. The refinements $\mathbf{GF}(\neg \mathit{hburst1})$ and $\mathbf{GF}\mathit{hlock0}$ (which are possible outputs of [3,26] if the user chooses the corresponding templates/variables) remove the first counterstrategy and refer to variables not in our interpolant. However neither is helpful.

8 Discussion

Targeting unrealizable cores The evaluation shows that our approach automatically selects variables that need to be constrained in order to reach realizability. In particular, all the intermediate refinements eliminate precisely a cause of unrealizability, consisting of the set of minimally unfulfillable guarantees from which a counterstrategy has been computed. We note that the returned variables in the AMBA02 example (*hready* and *hbusreq1*) are a subset of the variables that the authors in [3] suggest to use in order to instantiate the refinement templates.

Helpfulness of intermediate refinements is a desirable condition for reducing the convergence rate of the algorithm. When this holds, then the expected tree depth d (see Sect. 6) is reduced to $O(n_g)$, where n_g is the number of minimally unfulfillable subsets of guarantees. The application of our approach on the case studies consistently supports the attainment of this condition.

Number of unrollings We further define an upper bound to the number of unrollings needed to reach the termination condition as of Sect. 5.2. Every unrolling iteration produces an interpolant which is either the same as the previous iteration, or contains the description of one more state in the counterplay. In the worst case, without unrolling the interpolant describes just the initial state; after the first unrolling it contains a state component for the first transient state

s_1^{trans}; it is iteratively strengthened by one more state component until it describes all the transient states and the first replica of the unrolled states $s_{j,1}^{unr}$ for each j. In the following unrolling step, the interpolant contains the component of $s_{1,2}^{unr}$, which is such that $\lambda'(s_{1,2}^{unr}) = \lambda'(s_{1,1}^{unr})$: interpolation produces an equivalent refinement to the previous step, and therefore the procedure terminates. The maximum number of unrollings before reaching the termination condition is: $u_{C,MAX} = |S|$ where $|S|$ is the number of states in π_C.

Comparison with existing approaches Our approach extracts weaker refinements that those of [3]. The reason is that [3] uses templates that are true over all paths of a counterstrategy, while our approach requires them to be true in a single counterplay. More specifically, an invariant template used in [3] has the form $\mathbf{G}\neg q \vee \mathbf{X}\left(\bigwedge_{q' \in Next(q)} \neg q'\right)$, where q and q' indicate states in a counterstrategy and $Next(q)$ is the set of successor states of q; our approach extracts invariants of the form $\mathbf{G}\neg q \vee \neg q'$ for a $q' \in Next(q)$, which is implied by the former one, provided that they use the same variables set for q and q'.

We notice that in principle our approach may generate assumptions containing only output variables. This happens if some state component in the interpolant contains only output variables. Those are valid GR(1) formulae according to the definition, although hardly interpretable as constraints on the environment. Existing approaches circumvent the problem by allowing only input variables in their refinements [3,4]: however, in this way valid assumptions are also excluded. In our AMBA04 case study, one of the computed assumptions was $\mathbf{G}((\neg hmaster0 \wedge hbusreq1) \rightarrow \mathbf{X}(\neg hbusreq1))$, where $hmaster0 \in \mathcal{Y}$ and $hbusreq1 \in \mathcal{X}$. This assumption would not have been computed with that restriction.

9 Conclusions

We presented an interpolation-based approach for synthesizing weak environment assumptions for GR(1) specifications. Our approach exploits the information in counterstrategies and unrealizable cores to compute assumptions that directly target the cause of unrealizability. Compared to closely related approaches [3,26], our algorithm does not require the user to provide the set of variables upon which the assumptions are constructed. The case study applications show that our approach implicitly performs a variable selection that targets an unrealizable core, allowing for a quicker convergence to a realizable specification.

The final set of refinements is influenced by the choice of counterplay. We are investigating in our current work the effect of and criteria over the counterplay selection particularly on the full-separability of interpolants. Furthermore, since interpolants are over-approximations of the counterplays, the final specification is an under-approximation. In future work, we will explore the use of witnesses (winning strategies for the system) to counteract this effect. Finally, the applicability of our approach depends on the separability properties of the computed

interpolants: further investigation is needed to characterize the conditions under which an interpolation algorithm returns fully-separable interpolants.

Acknowledgments The support of the EPSRC HiPEDS CDT (EP/L016796/1) and Imperial College Junior Research Fellowship is gratefully acknowledged.

References

1. https://github.com/davidecavezza/TACAS17CaseStudies
2. Alrajeh, D., Kramer, J., van Lamsweerde, A., Russo, A., Uchitel, S.: Generating obstacle conditions for requirements completeness. In: ACM/IEEE 34th International Conference on Software Engineering. pp. 705–715 (2012)
3. Alur, R., Moarref, S., Topcu, U.: Counter-strategy guided refinement of GR(1) temporal logic specifications. In: Formal Methods in Computer-Aided Design. pp. 26–33 (2013)
4. Alur, R., Moarref, S., Topcu, U.: Pattern-Based Refinement of Assume-Guarantee Specifications in Reactive Synthesis. In: 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 501–516 (2015)
5. Angluin, D.: Learning regular sets from queries and counterexamples. Information and Computation 75(2), 87–106 (1987)
6. ARM Ltd.: AMBA-AHB specification. Available at https://developer.arm.com/products/architecture/amba-protocol
7. Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded Model Checking. Advances in Computers 58 (2003)
8. Bloem, R., Cimatti, A., Greimel, K., Hofferek, G., Könighofer, R., Roveri, M., Schuppan, V., Seeber, R.: RATSY A New Requirements Analysis Tool with Synthesis. In: 22nd International Conference on Computer Aided Verification, pp. 425–429 (2010)
9. Bloem, R., Jobstmann, B., Piterman, N., Pnueli, A., Sa'Ar, Y.: Synthesis of Reactive(1) designs. Journal of Computer and System Sciences 78(3), 911–938 (2012)
10. Bruttomesso, R., Cimatti, A., Franzén, A., Griggio, A., Sebastiani, R.: The MATHSAT 4 SMT solver: Tool paper. In: 20th International Conference on Computer Aided Verification. pp. 299–303 (2008)
11. Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Environment Assumptions for Synthesis. In: 9th International Conference on Concurrency Theory. pp. 147–161 (2008)
12. Cimatti, A., Roveri, M., Schuppan, V., Tchaltsev, A.: Diagnostic Information for Realizability. In: 9th International Conference on Verification, Model Checking, and Abstract Interpretation. pp. 52–67 (2008)
13. Cimatti, A., Griggio, A., Sebastiani, R.: Efficient Interpolant Generation in Satisfiability Modulo Theories. In: 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 397–412 (2008)
14. Cobleigh, J.M., Giannakopoulou, D., Păsăreanu, C.S.: Learning Assumptions for Compositional Verification. In: 9th International Conference Tools and Algorithms for the Construction and Analysis of Systems. pp. 331–346 (2003)
15. Craig, W.: Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. The Journal of Symbolic Logic 22(03), 269–285 (1957)
16. Degiovanni, R., Alrajeh, D., Aguirre, N., Uchitel, S.: Automated goal operationalisation based on interpolation and SAT solving. In: ACM/IEEE 36th International Conference on Software Engineering. pp. 129–139 (2014)
17. D'Silva, V., Kroening, D., Weissenbacher, G.: A Survey of Automated Techniques for Formal Software Verification. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 27(7), 1165–1178 (2008)
18. Esparza, J., Kiefer, S., Schwoon, S.: Abstraction Refinement with Craig Interpolation and Symbolic Pushdown Systems. In: 12th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, pp. 489–503 (2006)

19. Gheerbrant, A., ten Cate, B.: Craig Interpolation for Linear Temporal Languages. In: 23rd International Workshop on Computer Science Logic. pp. 287–301 (2009)
20. Gheorghiu Bobaru, M., Păsăreanu, C.S., Giannakopoulou, D.: Automated Assume-Guarantee Reasoning by Abstraction Refinement. In: 20th International Conference Computer Aided Verification. pp. 135–148 (2008)
21. Kamide, N.: Interpolation theorems for some variants of LTL. Reports on Mathematical Logic 50, 3–30 (2015)
22. Könighofer, R., Hofferek, G., Bloem, R.: Debugging formal specifications using simple counterstrategies. In: Formal Methods in Computer-Aided Design. pp. 152–159 (2009)
23. Krajíček, J.: Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. The Journal of Symbolic Logic 62(02), 457–486 (1997)
24. van Lamsweerde, A.: Requirements Engineering - From System Goals to UML Models to Software Specifications. Wiley (2009)
25. van Lamsweerde, A., Letier, E.: Handling obstacles in goal-oriented requirements engineering. IEEE Trans. Software Eng. 26(10), 978–1005 (2000)
26. Li, W., Dworkin, L., Seshia, S.A.: Mining assumptions for synthesis. In: ACM/IEEE 9th International Conference on Formal Methods and Models for Codesign. pp. 43–50 (2011)
27. Li, W., Sadigh, D., Sastry, S.S., Seshia, S.A.: Synthesis for human-in-the-loop control systems. In: 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 470–484 (2014)
28. Manna, Z., Pnueli, A.: The Temporal Logic of Reactive and Concurrent Systems. Springer (1992)
29. McMillan, K.L.: Interpolation and SAT-Based Model Checking. In: 15th International Conference on Computer Aided Verification. pp. 1–13 (2003)
30. Păsăreanu, C., Dwyer, M., Huth, M.: Assume-guarantee model checking of software: A comparative case study. Theoretical and Practical Aspects of SPIN Model Checking pp. 168–183 (1999)
31. Sohail, S., Somenzi, F.: Safety first: a two-stage algorithm for the synthesis of reactive systems. International Journal on Software Tools for Technology Transfer 15(5-6), 433–454 (2013)
32. Sohail, S., Somenzi, F., Ravi, K.: A Hybrid Algorithm for LTL Games. In: Verification, Model Checking, and Abstract Interpretation. pp. 309–323 (2008)

Appendix

Translation of GR(1) formulae into Boolean logic

Let $\phi = \phi^{\mathcal{E}} \rightarrow \phi^{\mathcal{S}}$ be a GR(1) property, expressed over a set $\mathcal{V}$ of propositional variables. Given a counterplay with states defined as in Section 5, the translation of ϕ^{θ} into Boolean logic is a Boolean (non-temporal) formula $[[\phi^{\theta}]]$ in the variables of the set $\mathcal{V}(S_{\pi})$ defined in Section 5.1. The three components of ϕ^{θ} are translated separately:

- An initial state condition $\varphi^{\theta}_{init} = B(\mathcal{V})$ is translated to $B(\mathcal{V}(s^{init}))$ by replacing each occurrence of $p \in \mathcal{V}$ with $p(s^{init}) \in \mathcal{V}(s^{init})$.
- An invariant $\varphi^{\theta}_{inv} = \mathbf{G}B(\mathcal{V} \cup \mathbf{X}\mathcal{V})$ is translated as the conjunction over all states in S_{π} $\bigwedge_{s \in S_{\pi}} B(\mathcal{V}(s) \cup \mathcal{V}(\mathrm{succ}\,(s)))$.
- Finally, a fairness condition $\varphi^{\theta}_{fair} = B(\mathcal{V})$ is translated as $\bigvee_{s \in S^{loop}} B(\mathcal{V}(s))$.

Proof of soundness [Theorem 1]

Proof. Since we are assuming that I_u is state-separable, and since by definition the counterplay translation implies I_u, each state component (which is a conjunct in I_u) is implied by $[[\pi_{C,u}, \varphi^{\mathcal{E}}_u]]$:

$$[[\pi_{C,u}, \varphi^{\mathcal{E}}_u]] \rightarrow I_u \rightarrow B_s(\mathcal{V}(s)) \tag{5}$$

for every s. By construction, a state component holds true iff it is satisfied by the corresponding state in the counterplay:

$$[[\pi_{C,u}, \varphi^{\mathcal{E}}_u]] \rightarrow B_s(\mathcal{V}(s)) \Leftrightarrow \langle \pi_C, s \rangle \models B_s(\mathcal{V}) \tag{6}$$

Now let us consider s^{init}. By (6) and I/O-separability:

$$\begin{aligned} [[\pi_{C,u}, \varphi^{\mathcal{E}}_u]] \rightarrow B^{init}(V(s^{init})) \Rightarrow & \langle \pi_C, s^{init} \rangle \models B^{init}(V) \\ \Rightarrow & \langle \pi_C, s^{init} \rangle \models B^{init}_{\mathcal{X}}(\mathcal{X}) \wedge B^{init}_{\mathcal{Y}}(\mathcal{Y}) \\ \Rightarrow & \pi_C \models B^{init}_{\mathcal{X}}(\mathcal{X}) \end{aligned} \tag{7}$$

So, π_C satisfies the part of the translation that refers to the initial state.

The next step is to consider pairs of consecutive states s and $\mathrm{succ}\,(s)$. Since I_u is I/O-separable, we have

$$B_{\mathrm{succ}(s)}(\mathcal{V}(\mathrm{succ}(s))) \rightarrow B_{\mathrm{succ}(s),\mathcal{X}}(\mathcal{X}(\mathrm{succ}(s)))$$

By (6) and LTL satisfaction definition for $\mathbf{X}$:

$$\begin{aligned} [[\pi_{C,u}, \varphi^{\mathcal{E}}_u]] \rightarrow B_{\mathrm{succ}(s),\mathcal{X}}(\mathcal{X}(\mathrm{succ}(s))) \Rightarrow & \langle \pi_C, \mathrm{succ}(s) \rangle \models B_{\mathrm{succ}(s),\mathcal{X}}(\mathcal{X}) \\ \Rightarrow & \langle \pi_C, s \rangle \models \mathbf{X}B_{\mathrm{succ}(s),\mathcal{X}}(\mathcal{X}) \\ \Rightarrow & \langle \pi_C, s \rangle \models B_{\mathrm{succ}(s),\mathcal{X}}(\mathbf{X}\mathcal{X}) \end{aligned} \tag{8}$$

From the conjunction of (6) and (8), and from the LTL interpretation of the operator **F** we finally get

$$\begin{aligned}[[\pi_{C,u}, \varphi_u^{\mathcal{E}}]] \rightarrow & B_s(\mathcal{V}(s)) \wedge B_{\mathrm{succ}(s)}(\mathcal{V}(\mathrm{succ}(s))) \\ \Rightarrow & \langle \pi_C, s \rangle \models B_s(\mathcal{V}) \wedge B_{\mathrm{succ}(s),\mathcal{X}}(\mathbf{X}\mathcal{X}) \\ \Rightarrow & \pi_C \models \mathbf{F}\left(B_s(\mathcal{V}) \wedge B_{\mathrm{succ}\,(s),\mathcal{X}}(\mathbf{X}\mathcal{X})\right)\end{aligned} \tag{9}$$

Therefore π_C satisfies the "eventually" subformulae in the translation.

Finally, let us consider the looping and unrolled states. In general, a path ending with a loop among states $q_1, \ldots, q_{|Q_{loop}|}$ satisfies the formula

$$\mathbf{FG} \bigvee_{j=1}^{|Q_{loop}|} B_{q_j}(\mathcal{V}) \tag{10}$$

where B_{q_j} is any Boolean expression that holds in q_j. The reason is that there is a suffix of the path that contains only states from Q_{loop}; therefore, this suffix always satisfies any of the looping states' valuations.

Now, by construction, unrolled states are replicates of looping states; therefore any Boolean formula that is true in an unrolled state is also true in the looping state of which it is a replicate:

$$\langle \pi_C, s_j^{loop} \rangle \models B_{j,r}^{unr}(\mathcal{V}) \Leftrightarrow \langle \pi_C, s_{j,r}^{unr} \rangle \models B_{j,r}^{unr}(\mathcal{V}) \tag{11}$$

Hence we have:

$$\langle \pi_C, s_j^{loop} \rangle \models \left(B_j^{loop}(\mathcal{V}) \wedge \bigwedge_{r=1}^{u} B_{j,r}^{unr}(\mathcal{V}) \right) \tag{12}$$

This can be replaced to the formula in (10) and we obtain:

$$\pi_C \models \mathbf{FG} \bigvee_{j=1}^{|S^{loop}|} \left(B_j^{loop}(\mathcal{V}) \wedge \bigwedge_{r=1}^{u} B_{j,r}^{unr}(\mathcal{V}) \right) \tag{13}$$

Therefore π_C satisfies also the conjuncts of $\mathcal{T}(I_u)$ that refer to looping states.

Since π_C satisfies each of the conjuncts in (3), then $\pi_C \models \mathcal{T}(I_u)$.

McMillan's Interpolation

Definition 5 (Unsatisfiability Proof [29]). *A proof of unsatisfiability for a set of clauses C is a directed acyclic graph (V, E), where the vertices V is a set of clauses, such that for every vertex $c \in V$, either*

- *$c \in C$, and c is a root, or*
- *c has exactly two predecessors, c_1 and c_2, and c is their resolvent,*

and the empty clause is the unique leaf.

Definition 6 (Interpolation Algorithm [29]).

Let ϕ_1 and ϕ_2 be a pair of clause sets and let $\Pi = (V, E)$ be a proof of unsatisfiability of $\phi_1 \wedge \phi_2$, with leaf vertex false. For all vertices $c \in V$, let $p(c)$ be a Boolean formula, such that

- *if c is a root, then*
  - *if $c \in \phi_1$ then $p(c) = g(c)$,*
  - *otherwise $p(c)$ is the constant true.*
- *else, considering c_1 and c_2 are the predecessors of c, v their pivot variable,*
  - *if v is in the language $\mathcal{L}_{\phi_1} \backslash \mathcal{L}_{\phi_2}$, then $p(c) = p(c1) \vee p(c2)$,*
  - *otherwise $p(c) = p(c1) \wedge p(c2)$.*